Bag of Attributes for Video Event Retrieval

Leonardo A. Duarte[1], Otávio A. B. Penatti[2], and Jurandy Almeida[1]

[1]Institute of Science and Technology
Federal University of São Paulo – UNIFESP
12247-014, São José dos Campos, SP – Brazil
Email: {leonardo.assuane, jurandy.almeida}@unifesp.br

[2]Advanced Technologies
SAMSUNG Research Institute
13097-160, Campinas, SP – Brazil
Email: o.penatti@samsung.com

arXiv:1607.05208v1 [cs.IR] 18 Jul 2016

***Abstract*—In this paper, we present the Bag-of-Attributes (BoA) model for video representation aiming at video event retrieval. The BoA model is based on a semantic feature space for representing videos, resulting in high-level video feature vectors. For creating a semantic space, i.e., the attribute space, we can train a classifier using a labeled image dataset, obtaining a classification model that can be understood as a high-level codebook. This model is used to map low-level frame vectors into high-level vectors (e.g., classifier probability scores). Then, we apply pooling operations on the frame vectors to create the final bag of attributes for the video. In the BoA representation, each dimension corresponds to one category (or attribute) of the semantic space. Other interesting properties are: compactness, flexibility regarding the classifier, and ability to encode multiple semantic concepts in a single video representation. Our experiments considered the semantic space created by a deep convolutional neural network (OverFeat) pre-trained on 1000 object categories of ImageNet. OverFeat was then used to classify each video frame and max pooling combined the frame vectors in the BoA representation for the video. Results using BoA outperformed the baselines with statistical significance in the task of video event retrieval using the EVVE dataset.**

I. Introduction

The retrieval of videos from specific events e.g., the wedding of Prince William and Kate Middleton, is a challenging application, as the goal is to retrieve other videos from that event in a database containing lots of different events. This task is even more challenging if we are considering only visual content, i.e., no textual annotations. Different events can occur at the same locations but in different dates, making videos of such events very similar visually. Other challenge is that there can be a large variation in visual aspects, even in the same event. For instance, for the wedding of Prince William and Kate Middleton, there can be videos with close-ups in the people and videos of the location (church, city buildings, etc).

Traditional video descriptors are usually based on low-level features, like textures and local patches [1]–[3], which rarely represent semantic properties. Some more recent approaches aim at at including semantics in image or video representations [4]–[6]. Object Bank [5], for instance, represents images by the responses to object filters. Bag of Scenes [4] considers a dictionary of scenes instead of a dictionary of local patches, and uses it for video representation aiming at geocoding, as the scenes can be representative of places.

If we could have a representation that can encode the multiple elements that appear in a given event in a single feature vector, we could better describe such event and discriminate it from others. Such a representation can be achieved by considering a classifier of high-level concepts in the video. Such concepts could be objects, scenes, locations, and so on.

To achieve such high-level representation for video event retrieval, we present the Bag-of-Attributes (BoA) model. The proposed model is based on a semantic feature space for representing video content, i.e., the attribute space. Such space can be created by training a classifier using a labeled dataset. Video contents can then be described by applying the learned classifier. The video vector contains the responses of the classifier, in which we have the activations of the semantic concepts that appears in the video. Such representation is a high-level feature vector for the video.

We validated the BoA model for video event retrieval using the EVVE dataset [7]. For obtaining the semantic feature space, we considered the use of OverFeat, a deep convolutional neural network pre-trained on 1000 object categories of ImageNet. The OverFeat model was used to classify each video frame and the probability scores were used as high-level feature vectors for the frames. The final video vector (the bag of attributes) was obtained by applying max pooling over the frame vectors. Results point that the BoA model outperform the baselines for video event retrieval.

II. Related Work

In this section, we describe related work specially devoted to video event retrieval, which is the focus of this paper.

The work of [7] shows a new approach for event retrieval in large datasets using only the visual properties of the videos, not considering audio neither textual metadata. Authors presented two methods and their combination and evaluated the results on EVVE dataset. Mean-MultiVLAD (MMV) is a method that averages frame descriptions extracted with SIFT and reduced by PCA. It aggregates the frame descriptions with VLAD,

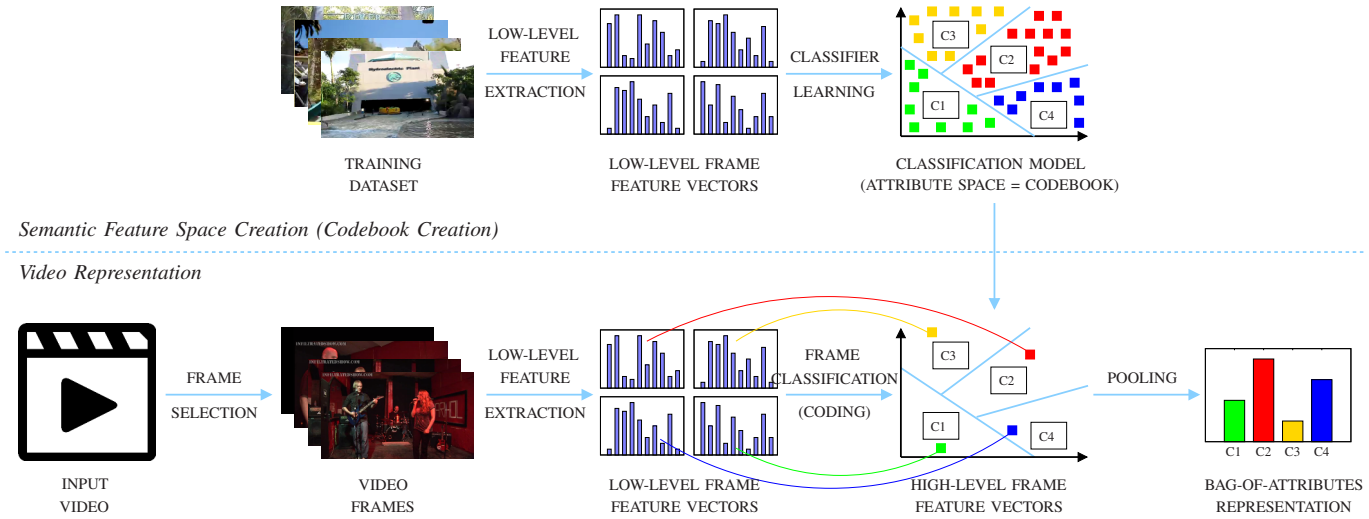

Fig. 1. Overview of the Bag-of-Attributes model. On top, we show how to obtain a semantic feature space for then using it to represent videos. At the bottom, we show how to map video content into this semantic space. As the process is similar to the creation and use of visual codebooks, we also show the names that are commonly used in that scenario. CN represent the categories (classes or attributes) in the training dataset.

generating the MMV representation. It does not include temporal information. They also propose a method to mix visual and temporal information in just one representation, called Circulant Temporal Encoding (CTE), which explores circulant matrices to enable a comparison in the frequency domain. On the experiments, they combined both methods creating the MMV+CTE representation inserting the normalized results obtained by the two methods for each video and for each query.

In the work of [8], authors improved the MMV method using a new hyper-pooling strategy to encode videos in a stable manner. They evaluated four different hashing functions: k-means, partial k-means (PKM), sign of stable components (SSC) e KD-Tree. The best result was obtained with SSC. Experiments also evaluated a representation made with Fisher Vectors as a baseline.

In [9], authors also show results on EVVE dataset, but the work is not comparable with ours, because all the experiments were evaluated in a classification task and the EVVE's official experimental procotol just includes a video event retrieval task. They also used some aditional datasets to train their methods on specific categories.

Methods related to the embedding of semantic features in the representation are also related to the proposed bag of attributes, although some of them are proposed for images [4]–[6], not for videos. Bag of Scenes (BoS) [4], originally proposed for video geocoding, uses visual codebooks based on whole scenes, instead of based on local patches (e.g., corners or edges), which represent places of interest. For geocoding, the BoS vector works as a place activation vector, helping in the geocoding task. The Object Bank (OB) [5] creates a semantic representation for *images* using a set of object filters. Then images are represented as the responses to such filters. In [6], authors proposed the SUN Attribute dataset and also a way of representing images according to an attribute classifier. OB and the method of [6] could be used as part of the proposed model, for representing each video frame in a semantic space.

III. Bag of Attributes

In this section, we present the Bag-of-Attributes (BoA) model for video representation. The main objective of the BoA model is to represent videos in a feature space with semantic information, resulting in a high-level representation [4], [5]. For that, we basically need to have a semantic feature space and a mapping function from the original video space to this new space. The steps involved in the BoA model are presented in Figure 1.

In the BoA model, we obtain the semantic feature space by training a supervised machine learning classifier based on a labeled dataset. The learned classifier, thus incorporates semantic information learned from the dataset. We call each label of the learning set as an *attribute*, aiming at referring to elements containing semantics.

For mapping or coding the video properties in this semantic (high-level) feature space, we start by classifying each frame of the input video with the learned classifier. Therefore, each frame is represented by its classification results, creating a high-level feature vector. Such results can be simply the class label given by the classifier or the whole probability vector, containing the probabilities of that frame in relation to every attribute of the learned classifier. Then, after having a high-level feature vector for each video frame, we generate the final video representation by computing some statistical measure over the frame vectors.

An obvious but important remark about the low-level feature extraction from video frames: in both stages (creation of the semantic space and video representation, i.e., top and bottom parts of Figure 1), the low-level feature space must be the same. For instance, if the classifier was trained with frames

TABLE I
EVVE EVENTS LIST. THE DATASET HAS A TOTAL OF 620 QUERY VIDEOS AND 2,375 DATABASE VIDEOS DIVIDED INTO 13 EVENTS. Q REFERS TO THE NUMBER OF QUERIES, DB+ AND DB- ARE THE NUMBERS OF POSITIVE AND NEGATIVE VIDEOS IN THE DATABASE, RESPECTIVELY.

ID	Event name	Q	Db+	Db-
1	Austerity riots in Barcelona, 2012	13	27	122
2	Concert of Die toten Hosen, Rock am Ring, 2012	32	64	143
3	Arrest of Dominique Strauss-Kahn	9	19	60
4	Egyptian revolution: Tahrir Square demonstrations	36	72	27
5	Concert of Johnny Hallyday stade de France, 2012	87	174	227
6	Wedding of Prince William and Kate Middleton	44	88	100
7	Bomb attack in the main square of Marrakech, 2011	4	10	100
8	Concert of Madonna in Rome, 2012	51	104	67
9	Presidential victory speech of Barack Obama 2008	14	29	56
10	Concert of Shakira in Kiev 2011	19	39	135
11	Eruption of Strokkur geyser in Iceland	215	431	67
12	Major autumn flood in Thailand, 2011	73	148	9
13	Jurassic Park ride in Universal Studios theme park	23	47	10
All	>>>	**620**	**1252**	**1123**

represented by color histograms, the classifier, of course, can only be applied over color histograms. Therefore, the frames of the video to be represented by BoA must have color histograms as low-level feature vectors.

We can easily map the steps in the BoA model to the steps involved in the context of visual codebooks and bags of visual words. The learned classifier can be seen as the *codebook*: each visual word is an attribute, i.e., a region in the classification space. The process of classifying each frame with the learned classifier can be seen as the *coding* step (visual word assignment). If in this step we consider only the classifier final attribution, i.e., class label for the frame, we have something similar to hard assignment. If we consider the classifier probability vector, we have something similar to soft assignment [10], [11]. Then, the final step of summarizing the frame representations can be seen as the *pooling* step, which can be implemented by summing, averaging or considering the maximum probability score among frames for each class [12].

Some interesting properties of the BoA representation are: (i) one dimension for each semantic concept, (ii) compactness (dimensionality equal to the number of classes in the learned classifier), (iii) flexibility to use any kind of classifier for creating the semantic feature space, and (iv) ability to encode multiple semantic concepts in a single representation. The last property can be understood if we consider that in the pooling operation we keep probability scores of the multiple classes activated along the video frames. For instance, if our attribute space is based on objects (like the object categories of ImageNet [13]), each frame will be classified considering the presence or not of such objects in the frame. The final video vector will then contain information of the objects that appear along the video. The BoA representation can be generalized to many other scenarios, which depend only on the attribute space to be considered. Other possible examples could be by considering classifiers trained to categorize scenes, faces, plants, vehicles, actions, etc.

For implementing the BoA model, different approaches can be used. Techniques like sampling at fixed-time intervals or summarization methods [14], [15] are examples of possibilities for frame selection. For creating the attribute classifier (i.e., the codebook), which is one of the key steps in the BoA model, one can learn the classifier in the dataset which better represents the contents of interest. Other option is to employ existing pre-trained classifiers, like the state-of-the-art classifiers based on convolutional neural networks [16]–[19]. In our experiments, we used OverFeat [17], which is a deep Convolutional Neural Network (ConvNet) trained on one thousand object categories of ImageNet dataset. In this case, as OverFeat integrates low-level feature extraction and classification, the low-level feature extraction step of the BoA model is implicit in OverFeat, i.e., we do not need to extract low-level feature vectors from video frames before applying OverFeat to classify them. For the coding and pooling steps, we considered the whole probability vector for each frame (similar to soft assignment) and then max pooling.

IV. EXPERIMENTS AND RESULTS

Experiments were conducted on the EVVE (EVent VidEo) dataset[1]: an event retrieval benchmark introduced by Revaud et al. [7]. The dataset is composed of 2,995 videos (166 hours) collected from YouTube[2]. Those videos are distributed among 13 event categories and are divided into a query set of 620 (20%) videos and a reference collection of 2,375 (80%) videos. Each event is treated as an independent subset containing some specific videos to be used as queries and the rest to be used as database for retrieval, as shown in Table I. It is a challenging benchmark since the events are localized in both time and space, for instance, the event 1 refers to the great riots and strikes that happened in the end of March 2012 at Barcelona, Spain, however, in the database, there are a lot of videos from different strikes and riots around the world.

EVVE uses a standard retrieval protocol: a query video is submitted to the system which returns a ranked list of similar videos. Then, we evaluate the average precision (AP) of each query and compute the mean average precision (mAP) per

[1]http://pascal.inrialpes.fr/data/evve/ (As of April, 2016).
[2]http://www.youtube.com (As of April, 2016).

event. The overall performance is assessed by the average of the mAPs (avg-mAP) obtained for all the events.

Our experiments followed the official experimental protocol created by [7]. Initially, each video in the dataset was represented by a BoA. With the BoA of each video, a given query video was used to retrieve the rest database videos, which were ranked according to the Euclidean (L2) distance between their BoAs. Finally, we used the dataset official tool to evaluate the retrieval results[3].

In this paper, video frames were selected using the well-known FFmpeg tool[4] in a sampling rate of one frame per second. The attribute classifier used is OverFeat[5] [17], which was the winner of two competitions of ImageNet 2013[6]. OverFeat was trained in 1,000 object classes of ImageNet. Therefore, each video frame classified by OverFeat has a feature vector of one thousand probability values, each corresponding to the probability of that frame having a specific object of ImageNet. With such representation, the BoA vector contains a summary of the objects present in the video. Notice that we analyze only the visual content, ignoring audio information and textual metadata.

We compared the BoA approach against three baselines [7]: Mean-MultiVLAD (MMV), CTE (Circulant Temporal Encoding) and a combination of both methods, known as MMV+CTE. Also, we considered the variations of MMV with the following hyper-pooling functions [8]: k-means, partial k-means (PKM), sign of stable componentes (SSC), KD-Tree and Fisher Vectors. To make a fair comparison, these approaches were selected with their best performance based on the results reported in [7], [8].

In Figure 2, we compare the BoA representation and the baseline methods with respect to the avg-mAP. As we can observe, the performance of the BoA representation outperformed all the baseline methods by a large margin.

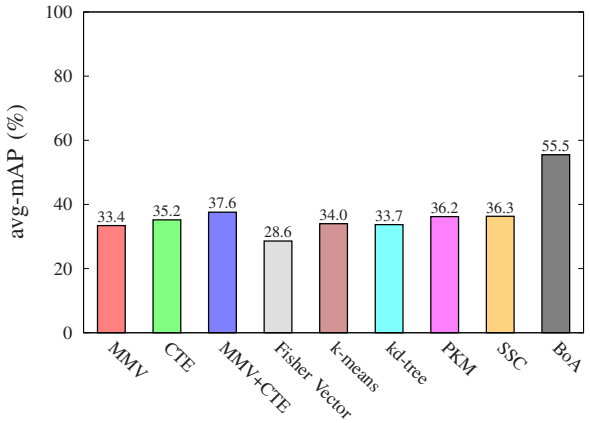

Fig. 2. Performance of different methods for event retrieval on EVVE dataset. The BoA representation outperformed all the baselines by a large margin.

[3]http://pascal.inrialpes.fr/data/evve/eval_evve.py (as of April 2016).
[4]http://www.ffmpeg.org/ (as of April 2016).
[5]http://cilvr.nyu.edu/doku.php?id=software:overfeat:start (as of April 2016).
[6]http://image-net.org/challenges/LSVRC/2013/ (as of April 2016).

The results where also compared by event, as shown in Table II. One can notice that BoA representation performed better than the baseline methods for most of the events (10 out of 13). For some events, the difference in favor of our method is very large, like in events 4, 5, 8, and 12.

TABLE II
RETRIEVAL PERFORMANCE (MAP) PER EVENT ON EVVE DATASET.

Event ID	MMV	CTE	MMV+CTE	BoA
1	23.90	13.90	24.60	**29.26**
2	19.90	16.60	20.20	**57.68**
3	8.70	12.80	11.10	**26.73**
4	12.60	10.80	13.20	**69.24**
5	23.40	26.20	26.00	**54.60**
6	33.80	41.30	39.40	**50.40**
7	12.40	**25.20**	21.20	13.86
8	25.40	25.70	28.10	**67.98**
9	53.10	**80.30**	69.40	43.65
10	45.50	40.90	**48.60**	33.87
11	77.30	71.40	77.40	**89.16**
12	36.60	29.70	37.10	**92.54**
13	60.40	69.30	71.90	**92.43**
avg-mAP	33.40	35.20	37.60	**55.49**

We also performed paired t-tests to verify the statistical significance of the results. For that, the confidence intervals for the differences between paired averages (mAP) of each category were computed to compare every pair of approaches. If the confidence interval includes zero, the difference is not significant at that confidence level. If the confidence interval does not include zero, then the sign of the difference indicates which alternative is better.

Figure 3 presents the confidence intervals (for $\alpha = 0.05$) of the differences between BoA and the baseline methods for the mAP measures. Notice that the confidence intervals for BoA and the baseline methods are always positive, indicating that BoA outperformed those approaches.

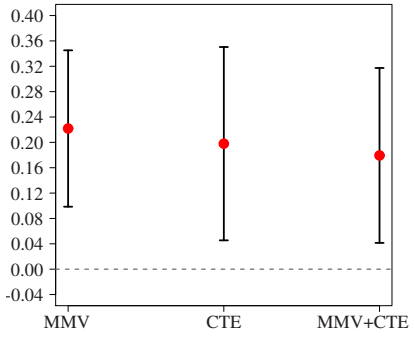

Fig. 3. Paired t-test comparing BoA and the baselines. As intervals are above zero, we can say that BoA outperformed the baselines with statistical significance.

According to the analysis of BoA results per event, one of the worst results happened on the event 1. On the other hand, one of the best results was obtained for the event 13. We perform a visual analysis of the videos to understand the main reasons for those results.

In case of the event 1 (see Figure 4), it is possible to see lots of riots and strikes at different places and moments. There are

(a) Positive videos (b) Negative videos

Fig. 4. Examples of video frames from Event 1 (Austerity riots in Barcelona, 2012), which was one of the events that BoA performed worst.

(a) Positive videos (b) Negative videos

Fig. 5. Examples of video frames from Event 13 (Jurassic Park ride in Universal Studios theme park), which was one of the events that BoA performed best.

(a) Positive videos (b) Negative videos

Fig. 6. Examples of video frames from Event 8 (Concert of Madonna in Rome, 2012), which was one of the events that BoA obtained the largest differences over the baselines.

scenes showing police, fire, cars, and crowd in almost all the videos (Figure 4b). Thus, it is difficult to identify only videos of the austerity riots that occurred in Barcelona at the end of March, 2012 (Figure 4a).

But, in case of the event 13 (see Figure 5), there are lots of similar positive videos, specially recorded at the entrance of

the ride, as shown in Figure 5a. This scene is repeated in many videos and probably helped our method. Negative videos do not contain the same entrance, as shown in Figure 5b.

Analyzing videos of the event 8 (see Figures 6), we can see that most positive videos are focusing on the main stage (Figure 6a). Therefore, there are lots of similar scenes. Negative videos include other Madonna concerts at different stages. Also, they usually show people, news and some other things related to the concert, but not just the main stage (Figure 6b).

We believe that our method outperformed the baselines because the proposed BoA representation carries semantic information. On the other hand, our method does not include temporal information and we think such feature is important to recognize some types of events.

V. Conclusions

We presented a new video representation for video event retrieval, named Bag of Attributes (BoA). In this model, videos are represented in a high-level feature space, which comprises the classification space defined by a supervised image classifier. In such space, each region corresponds to a semantic concept. To represent video content in this space, we start by classifying each video frame with the learned classifier, resulting in a high-level feature vector for each frame (e.g., classifier probability scores). Then, frame vectors are summarized by pooling operations to generate the final video vector, creating the BoA representation.

The main properties of the BoA representation are: each vector dimension corresponds to one semantic concept, compactness, flexibility regarding the learned classifier, and ability to encode multiple semantic concepts in a single vector.

In our implementation of the BoA model, we considered the semantic feature space created by a pre-trained deep convolutional neural network (OverFeat) trained on 1000 object categories of ImageNet. OveFeat was then used to classify each video frame and then max pooling was used to summarize the high-level frame vectors into the video bag of attributes. Our experiments considered the EVVE dataset for video event retrieval and the results showed that BoA outperformed the baselines with statistical significance. We believe that the ability to encode multiple concepts in the BoA representation could help discriminating between events.

As future work, we would like to evaluate other semantic spaces created by classifiers based on deep convolutional neural networks (e.g., AlexNet and GoogLeNet). We also would like to evaluate classifiers trained non-object categories, like scenes, for instance. The evaluation of the BoA model in other applications besides event retrieval is also a possible future work.

Acknowledgment

The authors would like to thank to CAPES, CNPq, and FAPESP for funding.

References

[1] J. Almeida, N. J. Leite, and R. S. Torres, "Comparison of video sequences with histograms of motion patterns," in *International Conference on Image Processing*, 2011, pp. 3673–3676.

[2] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu, "Action recognition by dense trajectories," in *Conference on Computer Vision and Pattern Recognition*, 2011, pp. 3169–3176.

[3] G. Willems, T. Tuytelaars, and L. Van Gool, "An efficient dense and scale-invariant spatio-temporal interest point detector," in *Computer Vision–ECCV*. Springer, 2008, pp. 650–663.

[4] O. A. B. Penatti, L. T. Li, J. Almeida, and R. da S. Torres, "A visual approach for video geocoding using bag-of-scenes," in *International Conference on Multimedia Retrieval*, 2012, pp. 1–8.

[5] L.-J. Li, H. Su, E. P. Xing, and F.-F. Li, "Object bank: A high-level image representation for scene classification and semantic feature sparsification," in *Neural Information Processing Systems*, 2010, pp. 1378–1386.

[6] G. Patterson, C. Xu, H. Su, and J. Hays, "The sun attribute database: Beyond categories for deeper scene understanding," *International Journal of Computer Vision*, vol. 108, no. 1, pp. 59–81, 2014.

[7] J. Revaud, M. Douze, C. Schmid, and H. Jégou, "Event retrieval in large video collections with circulant temporal encoding," in *Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2459–2466.

[8] M. Douze, J. Revaud, C. Schmid, and H. Jégou, "Stable hyper-pooling and query expansion for event detection," in *International Conference on Computer Vision*, 2013, pp. 1825–1832.

[9] X. Yang, T. Zhang, C. Xu, and M. S. Hossain, "Automatic visual concept learning for social event understanding," *IEEE Transactions on Multimedia*, vol. 17, no. 3, pp. 346–358, 2015.

[10] J. C. van Gemert, C. J. Veenman, A. W. M. Smeulders, and J.-M. Geusebroek, "Visual word ambiguity," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 7, pp. 1271–1283, 2010.

[11] L. Liu, L. Wang, and X. Liu, "In defense of soft-assignment coding," in *International Conference on Computer Vision*, 2011, pp. 1–8.

[12] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce, "Learning mid-level features for recognition," in *Conference on Computer Vision and Pattern Recognition*, 2010, pp. 2559–2566.

[13] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.

[14] J. Almeida, R. da S. Torres, and N. J. Leite, "Rapid video summarization on compressed video," in *IEEE International Symposium on Multimedia (ISM'10)*, 2010, pp. 113–120.

[15] J. Almeida, N. J. Leite, and R. da S. Torres, "VISON: VIdeo Summarization for ONline applications," *Pattern Recognition Letters*, vol. 33, no. 4, pp. 397–409, 2012.

[16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Neural Information Processing Systems*, 2012, pp. 1106–1114.

[17] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun, "Overfeat: Integrated recognition, localization and detection using convolutional networks," *CoRR*, vol. abs/1312.6229, 2013. [Online]. Available: http://arxiv.org/abs/1312.6229

[18] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Conference on Computer Vision and Pattern Recognition*, June 2015.

[19] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.